

A Mixed Observability Markov Decision Process Model for Musical Pitch

Pouyan Rafiei Fard

Department of Computer Engineering, Sharif University of Technology, Tehran, IRAN

RAFIEIFARD@CE.SHARIF.EDU

Keyvan Yahya

School of Psychology, University of Birmingham, Birmingham, UK

KXY054@BHAM.AC.UK

Abstract

Partially observable Markov decision processes have been widely used to provide models for real-world decision making problems. In this paper, we will provide a method in which a slightly different version of them called Mixed observability Markov decision process, MOMDP, is going to join with our problem. Basically, we aim at offering a behavioural model for interaction of intelligent agents with musical pitch environment and we will show that how MOMDP can shed some light on building up a decision making model for musical pitch conveniently.

1. Introduction

Partially observable Markov decision processes (POMDPs) have been widely used to provide models for real-world decision making problems. They provide a mathematical framework to model the interaction between the agent and its environment. One of the most notable characteristics of POMDPs is their ability to keep planning in dynamic environments and under uncertainty (Ong et al., 2010). To our knowledge, only a few authors have previously mentioned MDPs and POMDPs in the field of computer music. Among them, we could mention (Martin et al., 2010) who demonstrated the use of POMDPs to control musical behaviour in different conditions.

In this paper, we propose a novel model for interaction of the agents with musical pitch environment based on a variant of POMDPs called mixed observability Markov decision process (Ong et al., 2010).

In *5th International Workshop on Machine Learning and Music*, Edinburgh, Scotland, UK, 2012. Copyright 2012 by the author(s)/owner(s).

First, we mention the theoretical background of our work. In section 3, we propose our model for musical pitch based on MOMDPs. Section 4 addresses some implementation issues and presents an experiment to evaluate our model. Finally, we make our concluding remarks and discuss about the prospective potential developments of this models and its applications.

2. The Basic Idea of MOMDP

Beside the standard models of POMDP, there is a model called MOMDP that makes a slightly different with the former one. The latter is basically a factored POMDP which benefits from factorizing its states. In a MOMDP model, a state s is factored into two different variables x, y . So by writing $s = (x, y)$ we mean that s is consisted of two variables such that x stands for fully observable state and y stands for partially observable state. Thus having been factorized, we would have a mixed system space $S = X \times Y$ where X is the state of all values for x and either does Y for y .




Figure 1. the Standard POMDP model (left) and the MOMDP model (right) in which a state is divided into a fully-observable state x and a partially-observable state y (adapted from (Ong et al., 2010)).

3. The Proposed Model

From music theory, we know that any compound interval can be decomposed into some octaves and a simple interval which this idea can also be brought to any other simple intervals. In our MOMDP-based model, the agent makes its decisions according to the states that it receives from the environment which is here the musical pitch space. A MOMDP is denoted by the tuple $(X, Y, A, O, T_x, T_y, Z, R, \gamma)$. The relationship between these quantities and the musical concepts of our model are given elaborately in the following.

At each time step the environment is in a state $s \in S$ where $s = (x, y)$ and $x \in X$ is a fully observable state whereas $y \in Y$ is a partially observable one. In our model, a fully observable state represents a musical pitch in which the agent is having a precise estimate of its frequency at the time t plus the interval the agent is supposed to make. For the sake of simplicity, we only consider the natural musical pitches and the main intervals not beyond the octave interval. So, we have $S = \{C', D', E', F', G', A', B'\} \times \{1st', 2nd', 3rd', \dots, 7th'\}$. A is the set of actions available to the agent. Here an action $a \in A$ stands for a making a transition via a musical interval decomposition. In each state we define the possible actions with a set of decompositions. Relatively, the environment lies in partially observable states $y \in Y$ as the intermediate state regarding which one of actions the agent makes. Technically, the space of partially observable states is the same space for fully observable states. The parameter O is a set of observations that the agent makes which is the possible values of this parameter is the same as values from X and Y . Finally, The R parameter is $R_1\gamma + R_2$, where γ is the discount factor while R_1 is the reward for this first interval decomposition and R_2 is the reward given to the second one.

4. Experiment: Reinforcing patterns

For testing our model, we developed a Q-Learning algorithm (Watkins, 1989) to perform a similar task to which was done in (Cont, 2008). We made interactions with the system by feeding a relative pitch pattern as depicted in Figure 2., into the system. For this learning experiment, we set the learning parameters $\alpha = 0.4$, $\gamma = 0.5$ and $N = 20$ as the number of interactions. For a better demonstration of musical learning, the results are presented as intervals and notes. Thus, the y-axis of Figure 3. indicates the intervals and the x-axis is for the notes and for each interval-note pair. The gray-scale values show the learned Q-value and the intensity of these them shows the policy learned

by the agent. Also, the values indicated with red rectangles are the values which was originally fed into the system via the pitch contour.



Figure 2. Pitch contour pattern used in the experiment.




Figure 3. The results of the experiment.

5. Conclusion

The results of our experiments imply that our agent efficiently learned a behaviour policy. In addition, from Figure 3. we can see that not only our agent learned the given pitch contour (shown by red rectangles) but also some other state-action pairs. This is mainly happened because our method benefits from factorizing each state into a couple of fully-observable and partially-observable states. So, this approach will obviously help to have a faster convergence of the agent which is interacting with musical pitch environment.

References

- Ong, S. C. W., Png, S. W., Hsu D., and Lee, W. S. Planning under Uncertainty for Robotic Tasks with Mixed Observability. *International Journal of Robotics Research*, 29(8):1053–1068, 2010.
- Martin, A., Jin, C., van Schaik, A., and Martens, W. L. Partially Observable Markov Decision Processes for Interactive Music Systems. In *Proceedings of the International Computer Music Conference (ICMC 2010)*, pp. 490–493, New York, 2010.
- Watkins, C. J. C. H. *Learning from Delayed Rewards*. PhD thesis, Cambridge University, 1989.
- Cont, A. *Modeling Musical Anticipation: From the time of music to the music of time*. PhD thesis, University of Paris 6 and University of California in San Diego, 2008.